arXiv:1603.04118v2 [stat.ML] 22 Jun 2016

Active Algorithms For Preference Learning Problems with Multiple Populations

Aniruddha Bhargava*[1], Ravi Sastry Ganti†[2], and Robert Nowak‡[1]

[1]University of Wisconsin, Madison WI
[2]Walmart Labs, San Bruno, CA

June 23, 2016

Abstract

In this paper we model the problem of learning preferences of a population as an active learning problem. We propose an algorithm can adaptively choose pairs of items to show to users coming from a heterogeneous population, and use the obtained reward to decide which pair of items to show next. We provide computationally efficient algorithms with provable sample complexity guarantees for this problem in both the noiseless and noisy cases. In the process of establishing sample complexity guarantees for our algorithms, we establish new results using a Nyström-like method which can be of independent interest. We supplement our theoretical results with experimental comparisons.

1 Introduction

In this work, we are interested in designing active learning algorithms for preference learning problems where there are multiple sub-populations of users, and users from the same sub-population have similar preferences over a large set of items. We propose models where the algorithm chooses a pair of items for each incoming user and obtains a scalar reward which is possibly stochastic. The reward is large when either of the chosen pair of items is liked by the user. Our goal is to find an optimal pair of items by actively eliciting responses from users on different pairs of items.

As a concrete example, suppose an advertising engine can show just two advertisements to each incoming user. Different ads may appeal more or less to different sub-populations, but the right combination of two ads optimizes the chance that a random user will like one or the other. This paper focuses on selecting a pair of items, rather than larger subsets, for the sake of simplicity. We will discuss generalizations and extensions later in the paper.

Preference modeling problems such as the one described above and more are ubiquitous and show up in internet advertising [1], econometrics [2], and marketing [3]. An informal description of the preference learning problem that we consider in this paper is as follows: We are given a very large population of users and a set of K items. In most of the applications that we are interested in K is large. Suppose each user belongs to one of r unknown sub-populations of the original population. We are interested in learning from user data which pair of items, among the set of K items are maximally preferred. The contributions of this paper can be summarized as follows

1. We introduce two models for preference the modeling problem, when there are r sub-populations and K items. In each round we choose a pair of items and receive a reward which is a function of the pair chosen. This reward is large if either of the items in the chosen pair is "good". For both the models we are interested in

*aniruddha@wisc.edu
†gmravi2003@gmail.com.
‡nowak@ece.wisc.edu
The first two authors made equal contributions to the paper.
Most of the work was done while RSG was at UW-Madison.

designing algorithms that discover an (ϵ, δ) best pair of items using as few trials as possible i.e. algorithms which can output, with probability at least $1 - \delta$, a pair of items that is ϵ close to the best pair of items in terms of the expected reward of the pair. The difference between the models is whether the reward is stochastic or deterministic.

2. *The core idea behind our proposed algorithms is to reduce the problem of finding a near-optimal pair of items to the problem of finding a near-smallest element of a certain low-rank, symmetric positive semi-definite matrix (SPSD).* The rank of the matrix is equal to the number of sub-populations. While one could, in principle, use low-rank matrix completion (LRMC) techniques as in [4], our proposed algorithms explicitly exploit the SPSD structure in the problem.

3. The algorithms that we propose for both the models are based on iteratively finding linearly independent columns of an SPSD matrix. The key challenge in designing this column selection step is that querying elements of the SPSD matrix is either very expensive, or noisy, and hence we need to be selective about what elements of the matrix to query.

4. For all of the algorithms we establish sample complexity bounds of finding an (ϵ, δ) optimal pair of items. Establishing such sample complexity bounds leads to interesting problems in matrix approximation in the max-norm. The contributions we make here could be of independent interest in the low-rank matrix approximation literature.

To the best of our knowledge our formulation of preference learning problems as bandit problems with a large number of items and identifying the matrix structure in the items and deriving bounds in terms of rank of the underlying matrix is new and is what makes this work novel compared to other relevant literature.

Notation. Δ_r represents the r dimensional probability simplex. Matrices and vectors are represented in bold font. For a matrix $\boldsymbol{L}$, unless otherwise stated, the notation $\boldsymbol{L}_{i,j}$ represents (i, j) element of $\boldsymbol{L}$, and $\boldsymbol{L}_{i:j,k:l}$ is the submatrix consisting of rows $i, i+1, \ldots, j$ and columns $k, k+1, \ldots, l$. The matrix $\|\cdot\|_1$ and $\|\cdot\|_2$ norms are always operator norms. The matrix $\|\cdot\|_{\max}$ is the element wise infinity norm. Finally, let $\mathbb{1}$ be the all 1 column vector.

2 Related Work

Our framework is closely tied to literature on pure exploration problems in multi-armed bandits. In such pure exploration problems one is interested in designing algorithms with low simple regret or designing algorithms with low (ϵ, δ) query complexity. Algorithms with small simple regret have been designed in the past [5, 6, 7]. [8] suggested the Successive Elimination (SE) and Median Elimination (ME) to find near optimal arms with provable sample complexity guarantees. These sample complexity guarantees typically scale linearly with the number of arms. In principal, one could naively reduce our problem to a pure exploration problem where we need to find an (ϵ, δ) good arm. However, such naive reductions throw away any dependency information among the arms. The algorithms that we give build on the SE algorithm but crucially exploits the matrix structure in the problem to give much better algorithms than a naive reduction.

Bandit problems where multiple actions are selected have also been considered in the past. [9] consider a setup where on choosing multiple arms the reward obtained is the sum of the rewards of the chosen arms, and the reward of each chosen arm is reveled to the algorithm. Both these works focus on obtaining guarantees on the cumulative regret compared to the best set of arms in hindsight. [10] consider a problem, in the context of information retrieval, where multiple bandit arms are chosen and the reward obtained is the maximum of the rewards corresponding to the chosen arms. Apart from this reward information the algorithm also gets a feedback which is which one of the chosen arms has the highest reward. [11, 12] also study similar models. A major difference between the above mentioned works and our work is the feedback and reward model and the fact that we are not interested in regret guarantees but rather in finding a good pair of arms as quickly as possible. Furthermore our linear-algebraic approach to the problem is very different from previous approaches which were either based on multiplicative weights [9] or online greedy submodular maximization [11, 12, 10]. [13] also consider similar subset selection problems and provide algorithms to identify the top set of arms. In the Web search literature click models have been proposed to model user behaviour [14, 15]

and a bandit analysis of such models have also been proposed [16]. However, these models assume that all the users come from a single population and tend to use richer information in their formulations (for example information about which exact link was clicked). Interactive collaborative filtering (CF) and bandit approaches to such problems have also been investigated[17]. Though, the end goal in CF is different from our goal in this paper.

In our paper, we reduce the problem of preference learning to a low-rank SPSD matrix completion problem both in the noiseless and noisy cases. Many other authors have studied the problem of SPSD matrix completion [18]. However, all of these papers consider the passive case, i.e. the entries of the matrix that have been revealed are not under their control. In contrast, we have an active setup, where we can decide which entries in the matrix to reveal. The Nyström algorithm for approximation of low rank SPSD matrices has been well studied both empirically and theoretically. Nyström methods typically choose random columns to approximate the original low-rank matrix [19, 20]. Adaptive schemes where the columns used for Nystrom approximation are chosen adaptively have also been considered in the literature. To the best of the knowledge these algorithms either need the knowledge of the full matrix [21] or have no provable theoretical guarantees [22]. Moreover, to the best of our knowledge all analysis of Nystrom approximation that has appeared in the literature assume that one can get error free values for entries in the matrix which does not apply to the problems that we are interested in.

3 Models for preference learning problem

In this section, we propose two models for the preference modeling problem. In both the proposed models the algorithm chooses a pair of items, and receives a reward which is a function of the the chosen pair and possibly other external random factors. For both the models we want to design algorithms that discover, using as few trials as possible, an (ϵ, δ) best pair of items. The difference between the two models is that in the first model the reward obtained is a deterministic function of the pair of items chosen, whereas in our second model the reward obtained is random.

Model 3.1 Description of our proposed models

1: **while** TRUE **do**
2: In the case of stochastic model, nature chooses $Z_t \sim \mathrm{Mult}(\boldsymbol{p})$, but does not reveal it to the algorithm.
3: Algorithm chooses a pair of items (i_t, j_t).
4: Algorithm receives the reward y_t defined as follows

$$y_{t,\text{det}} = 1 - \mathbb{E}_{Z_t \sim \boldsymbol{p}}(1 - \boldsymbol{u}_{Z_t}(i_t))(1 - \boldsymbol{u}_{Z_t}(j_t)) \text{ // if model is deterministic} \tag{1}$$

$$y_{t,\text{stoc}} = \max\{y_{i_t}, y_{j_t}\} \text{ // if model is stochastic} \tag{2}$$

$$y_{i_t} \sim \mathrm{Bern}(\boldsymbol{u}_{Z_t}(i_t)) \tag{3}$$

$$y_{j_t} \sim \mathrm{Bern}(\boldsymbol{u}_{Z_t}(j_t)) \tag{4}$$

5: Algorithm stops if it has found a certifiable (ϵ, δ) optimal pair of items.
6: **end while**

In Figure (3.1) we sketch both the deterministic and the stochastic model. Let, Z_t be a multinomial random variable defined by a probability vector $\boldsymbol{p} \in \Delta_r$, whose output space is the set $\{1, 2, \ldots, r\}$. Let $\boldsymbol{u}_{Z_t}$ be a reward vector in $[0, 1]^K$ indexed by Z_t. On playing the pair of arms (i_t, j_t) in round t the algorithm receives a scalar reward y_t. In the deterministic model y_t is deterministic and is given by Equation (1), whereas in the stochastic model y_t is a random variable that depends on the random variable Z_t as well as additional external randomness. However, a common aspect of both these models is that the expected reward associated with the pair of choices (i_t, j_t) in round t is the same and is equal to the expression in given in Equation (1)

As a concrete example of the above models consider the example that we discussed in the introduction. Here, we have a random customer that belongs to one of the three possible categories with probability defined by a probability vector $\boldsymbol{p} \in \Delta_3$. When using a stochastic model, this advertising company which has no information about which subpopulation this customer is from, makes \$1 whenever, this customer clicks on either of the one of the two displayed

advertisements, and \$0 otherwise. This reward depends on two sources of randomness. (i) The sub-population that the user comes from which is modeled by a multinomial random variable Z_t and (ii) external randomness that is modeled by Bernoulli random variables with parameters $\boldsymbol{u}_{Z_t}(i_t)$ and $\boldsymbol{u}_{Z_t}(j_t)$. In the case of a deterministic model the reward obtained is the same as the expected reward in stochastic model, and therefore not random. It is clear from Figure (3.1) that the optimal pair of choices satisfies the equation

$$(i_\star, j_\star) = \arg\min_{i,j} \mathbb{E}_{Z_t \sim \boldsymbol{p}}(1 - \boldsymbol{u}_{Z_t}(i))(1 - \boldsymbol{u}_{Z_t}(j)). \tag{5}$$

Since, we are interested in returning an (ϵ, δ) optimal pair of choices it is enough if the pair returned by our algorithm attains an objective function value that is at most ϵ more than the optimal value of the objective function shown in equation (5), with probability at least $1 - \delta$.

4 Low rank structure in our models

Let $\boldsymbol{p} \in \Delta_r$, and let the reward matrix $\boldsymbol{R} \in \mathbb{R}^{K \times K}$ be such that its $(i,j)^{\text{th}}$ entry is the probability of obtaining a reward of 1 when the pair of items (i, j) are pulled. Then from equation (1) and equation (2) we know that the reward structure for both the deterministic and stochastic models has the form

$$\begin{aligned} \boldsymbol{R}_{i,j} &= 1 - \mathbb{E}_{Z_t \sim \boldsymbol{p}}(1 - \boldsymbol{u}_{Z_j}(i))(1 - \boldsymbol{u}_{Z_j}(j)) \\ &= 1 - \sum_{k=1}^{r} \boldsymbol{p}_k (1 - \boldsymbol{u}_k(i))(1 - \boldsymbol{u}_k(j)) \end{aligned} \tag{6}$$

$$\boldsymbol{R} = \mathbb{1}\mathbb{1}^\top - \underbrace{\sum_{k=1}^{r} \boldsymbol{p}_k (\mathbb{1} - \boldsymbol{u}_k)(\mathbb{1} - \boldsymbol{u}_k)^\top}_{\boldsymbol{L}}. \tag{7}$$

As mentioned before our goal is to find a pair of items that are (ϵ, δ) optimal. Hence, it is enough to find an entry in the matrix $\boldsymbol{L}$ that is ϵ close to the smallest entry in the matrix $\boldsymbol{L}$ with probability at least $1 - \delta$. A naive way to solve this problem is to treat this problem as a best-arm identification problem in stochastic multi-armed bandits where there are $\Theta(K^2)$ arms each corresponding to a pair of items. One could now run a Successive Elimination (SE) algorithm or a Median Elimination algorithm on these $\Theta(K^2)$ pairs [8] to find an (ϵ, δ) optimal pair. The sample complexity of the SE or ME algorithms on these $\Theta(K^2)$ pairs would be roughly $\tilde{\mathcal{O}}(\frac{K^2}{\epsilon^2})$ [1]. In typical applications that we are interested in, K can be very large, and therefore the sample complexity of such naive algorithms can be very large. However, these simple reductions throw away information between different pairs of items and hence are sub-optimal. A natural question to ask is can we design algorithms that can efficiently exploit the matrix structure in our problem? It turns out that in our problem we can get away with sample complexity far smaller than K^2. In order to do this we exploit the structural properties of matrix $\boldsymbol{L}$. From equation (7) it is clear that the matrix $\boldsymbol{L}$ can be written as a sum of r rank-1 matrices. Hence $\text{rank}(\boldsymbol{L}) \leq r$. Furthermore, since these rank-1 matrices are all positive semi-definite and $\boldsymbol{L}$ is a convex combination of such, we can conclude that $\boldsymbol{L} \succeq 0$. We exploit these properties in our algorithm design. We have proved the following proposition:

Proposition 4.1. *The matrix $\boldsymbol{L}$ shown in equation* (7) *satisfies the following two properties: (i)* $\text{rank}(\boldsymbol{L}) \leq r$ *(ii)* $\boldsymbol{L} \succeq 0$.

5 Algorithms for deterministic model

Our approach to finding an (ϵ, δ) pair in the deterministic model considered in this paper is via matrix completion. Choosing the pair of items (i, j) reveals the $(i,j)^{th}$ entry of the matrix $\boldsymbol{R}$. As can be seen from equation (7) the matrix $\boldsymbol{R}$ has rank at most $r + 1$. One can use standard low rank matrix completion (LRMC) ideas which relies

[1] The $\tilde{\mathcal{O}}$ notation hides logarithmic dependence on $\frac{1}{\delta}$, K, $\frac{1}{\delta}$

on nuclear norm minimization techniques [4]. A typical result from the LRMC literature says that we need to see $\mathcal{O}(Kr\mu\log^2(K))$ random entries in the matrix $\boldsymbol{R}$, where μ is the upper bound on the coherence of the row and column space of matrix $\boldsymbol{R}$, for exact recovery of the matrix $\boldsymbol{R}$. We provide a simple algorithm that recovers the matrix $\boldsymbol{R}$ after querying $\Theta(Kr)$ entries in the matrix $\boldsymbol{R}$. Our algorithm called PLANS [2] is shown in Figure (5.1) and works with the matrix $\boldsymbol{L}$ [3].

PLANS is shown in Figure (5.1). It is an iterative algorithm that finds out what columns of the matrix are independent. PLANS maintains a set of indices (denoted as $\mathcal{C}$ in the pseudo-code) corresponding to independent columns of matrix $\boldsymbol{L}$. Initially $\mathcal{C} = \{1\}$. PLANS then makes a single pass over the columns in $\boldsymbol{L}$ and checks if the current column is independent of the columns in $\boldsymbol{C}$. This check is done in line 5 of Figure (5.1) and most importantly requires *only the principal sub-matrix*, of $\boldsymbol{L}$, indexed by the set $\mathcal{C} \cup \{c\}$. If the column passes this test then all the elements in this column i whose values have not been queried in the past are queried and the matrix $\hat{\boldsymbol{L}}$ is updated with these values. The test in line 5 is the column selection step of the PLANS algorithm and is justified by Proposition (5.1). Finally, once r independent columns have been chosen, we impute the matrix by using Nystrom extension. Nystrom based methods have been proposed in the past to handle large scale kernel matrices in the kernel based learning literature [20, 22]. The major difference between this work and ours is that the column selection procedure in our algorithms is deterministic, whereas in Nystrom methods columns are chosen at random. The following proposition simply follows from the fact that any principal submatrix of an SPSD matrix is also SPSD and hence admits an eigen-decomposition.

Algorithm 5.1 Preference Learning via Adaptive Nystrom Sampling (PLANS)

Input: A deterministic oracle that takes a pair of indices (i, j) and outputs $\boldsymbol{L}_{i,j}$.

Output: $\hat{\boldsymbol{L}}$

1: Choose the pairs $(j, 1)$ for $j = 1, 2, \ldots, K$ and set $\hat{\boldsymbol{L}}_{j,1} = \boldsymbol{L}_{j,1}$. Also set $\hat{\boldsymbol{L}}_{1,j} = \boldsymbol{L}_{j,1}$
2: $\mathcal{C} = \{1\}$ {Set of independent columns discovered till now}
3: **for** $(c = 2;\ c \leftarrow c + 1;\ c \leq K)$ **do**
4: Query the oracle for (c, c) and set $\hat{\boldsymbol{L}}_{c,c} \leftarrow \boldsymbol{L}_{c,c}$
5: **if** $\sigma_{\min}\left(\hat{\boldsymbol{L}}_{\mathcal{C}\cup\{c\},\mathcal{C}\cup\{c\}}\right) > 0$ **then**
6: $\mathcal{C} \leftarrow \mathcal{C} \cup \{c\}$
7: Query $\mathcal{O}$ for the pairs $(\cdot, c)$ and set $\hat{\boldsymbol{L}}(\cdot, c) \leftarrow \boldsymbol{L}(\cdot, c)$ and by symmetry $\hat{\boldsymbol{L}}(c, \cdot) \leftarrow \boldsymbol{L}(\cdot, c)$.
8: **end if**
9: **if** $(|\mathcal{C}| = r)$ **then**
10: break
11: **end if**
12: **end for**
13: Suppose the submatrix of $\boldsymbol{L}$ corresponding to columns in $\mathcal{C}$ is $\boldsymbol{C}$ and the principal submatrix in $\boldsymbol{L}$ corresponding to indices in $\mathcal{C}$ is $\boldsymbol{W}$. Then, construct the Nystrom extension $\widehat{\boldsymbol{L}} = \boldsymbol{C}\boldsymbol{W}^{-1}\boldsymbol{C}^\top$.

Proposition 5.1. *Let $\boldsymbol{L}$ be any SPSD matrix of size K. Given a subset $\mathcal{C} \subset \{1, 2, \ldots, K\}$, the columns of the matrix $\boldsymbol{L}$ indexed by the set $\mathcal{C}$ are independent iff the principal submatrix $\boldsymbol{L}_{\mathcal{C},\mathcal{C}}$ is non-degenerate, equivalently iff, $\lambda_{\min}(\boldsymbol{L}_{\mathcal{C},\mathcal{C}}) > 0$.*

It is not hard to verify the following theorem. The proof has been relegated to the appendix.

Theorem 5.2. *If $\boldsymbol{L} \in \mathbb{R}^{K\times K}$ is an SPSD matrix of rank r, then the matrix $\hat{\boldsymbol{L}}$ output by the PLANS algorithm* (5.1) *satisfies $\hat{\boldsymbol{L}} = \boldsymbol{L}$. Moreover, the number of oracle calls made by PLANS is at most $K(r + 1)$. The sampling algorithm* (5.1) *requires: $K + (K - 1) + (K - 2) + \ldots + (K - (r - 1)) + (K - r) \leq (r + 1)K$ samples from the matrix $\boldsymbol{L}$.*

The following corollary follows immediately.

[2] Preference Learning via Adaptive Nystrom Sampling

[3] Pulling the pair (i,j) gets us $\boldsymbol{R}_{i,j}$, which can be used to populate $\boldsymbol{L}_{i,j} = 1 - \boldsymbol{R}_{i,j}$

Corollary 5.3. *Using algorithm* (5.1) *we can output a* $(0,0)$ *optimal pair of items for the deterministic model by obtaining rewards of at the most* $K(r+1)$ *pairs of items adaptively.*

Note that the sample complexity of the PLANS algorithm is slightly better than typical sample complexity results for LRMC. We managed to avoid factors logarithmic in dimension and rank, as well as incoherence factors that are typically found in LRMC results [4]. Also, our algorithm is purely deterministic, whereas LRMC uses randomly drawn samples from a matrix. In fact, this careful, deterministic choice of entries of the matrix is what helps us do better than LRMC.

Moreover, PLANS algorithm is optimal in a min-max sense. This is because any SPSD matrix of size K and rank r is characterized via its eigen decomposition by Kr degrees of freedom. Hence, any algorithm for completion of an SPSD matrix would need to see at least Kr entries. As shown in theorem (5.2) the PLANS algorithm makes at most $K(r+1)$ queries and hence PLANS is min-max optimal.

The PLANS algorithm needs the rank r as an input. However, the PLANS algorithm can be made to work even if r is unknown by simply removing the condition on line 9 in the PLANS algorithm. Even in this case the sample complexity guarantees in Theorem (5.2) hold. Finally, if the matrix is not exactly rank r but can be approximated by a matrix of rank r, then PLANS can be made more robust by modifying line 5 to $\sigma_{\min}\left(\hat{\boldsymbol{L}}_{\mathcal{C}\cup\{c\},\mathcal{C}\cup\{c\}}\right) \geq \sigma_{\text{thresh}}$, where σ_{thresh} depends on r.

6 Algorithms for stochastic model

For the stochastic model considered in this paper we shall propose an algorithm, called R-PLANS [4], which is a robust version of PLANS. Like PLANS, the robust version discovers a set of independent columns iteratively and then uses the Nyström extension to impute the matrix. Figure (6.1) provides a pseudo-code of the R-PLANS algorithm.

R-PLANS like the MISA algorithm repeatedly performs column selection steps to select a column of the matrix $\boldsymbol{L}$ that is linearly independent of the previously selected columns, and then uses these selected columns to impute the matrix via a Nystrom extension. In the case of deterministic models, due to the presence of a deterministic oracle, the column selection step is pretty straight-forward and requires calculating the smallest singular-value of certain principal sub-matrices. In contrast, for stochastic models the availability of a much weaker stochastic oracle makes the column selection step much harder. We resort to the successive elimination algorithm where principal sub-matrices are repeatedly sampled to estimate the smallest singular-values for those matrices. The principal sub-matrix that has the largest smallest singular-value determines which column is selected in the column selection step.

Given a set $\mathcal{C}$, define $\boldsymbol{C}$ to be a $K \times r$ matrix corresponding to the columns of $\boldsymbol{L}$ indexed by $\mathcal{C}$ and define $\boldsymbol{W}$ to be the $r \times r$ principal submatrix of $\boldsymbol{L}$ corresponding to indices in $\mathcal{C}$. R-PLANS constructs estimators $\widehat{\boldsymbol{C}}, \widehat{\boldsymbol{W}}$ of $\boldsymbol{C}, \boldsymbol{W}$ respectively by repeatedly sampling independent entries of $\boldsymbol{C}, \boldsymbol{W}$ for each index and averaging these entries. The sampling is such that each entry of the matrix $\boldsymbol{C}$ is sampled at least m_1 times and each entry of the matrix $\boldsymbol{W}$ is sampled at least m_2 times, where

$$m_1 = 100C_1(W,C)\log(2Kr/\delta)\max\left(\frac{r^{5/2}}{\epsilon}, \frac{r^2}{\epsilon^2}\right), m_2 = 200C_2(W,C)\log(2r/\delta)\max\left(\frac{r^3}{\epsilon}, \frac{r^5}{\epsilon^2}\right) \tag{8}$$

and C_1, C_2 are problem dependent constants defined as

$$C_1(\boldsymbol{W},\boldsymbol{C}) = \max\big(\left\|\boldsymbol{W}^{-1}\boldsymbol{C}^\top\right\|_{\max}, \left\|\boldsymbol{W}^{-1}\boldsymbol{C}^\top\right\|_{\max}^2, \left\|\boldsymbol{W}^{-1}\right\|_{\max}, \left\|\boldsymbol{C}\boldsymbol{W}^{-1}\right\|_1^2, \left\|\boldsymbol{W}^{-1}\right\|_2 \left\|\boldsymbol{W}^{-1}\right\|_{\max}\big) \tag{9}$$

$$C_2(\boldsymbol{W},\boldsymbol{C}) = \max\left(\left\|\boldsymbol{W}^{-1}\right\|_2^2 \left\|\boldsymbol{W}^{-1}\right\|_{\max}^2, \left\|\boldsymbol{W}^{-1}\right\|_2 \left\|\boldsymbol{W}^{-1}\right\|_{\max}, \left\|\boldsymbol{W}^{-1}\right\|_2, \left\|\boldsymbol{W}^{-1}\right\|_2^2\right) \tag{10}$$

R-PLANS then returns the Nyström extension constructed using matrices $\widehat{\boldsymbol{C}}, \widehat{\boldsymbol{W}}$.

[4] R-PLANS stands for Robust Preference Learning via Adaptive Nyström Sampling

Algorithm 6.1 Robust Preference Learning via Adaptive Nystrom Sampling (R-PLANS)

Input: $\epsilon > 0, \delta > 0$ and a stochastic oracle $\mathcal{O}$ that when queried with indices (i, j) outputs a Bernoulli random variable $\mathrm{Bern}(L_{i,j})$

Output: A SPSD matrix $\hat{\boldsymbol{L}}$, which is an approximation to the unknown matrix $\boldsymbol{L}$, such that with probability at least $1 - \delta$, all the elements of $\hat{\boldsymbol{L}}$ are within ϵ of the elements of $\boldsymbol{L}$.

$\mathcal{C} \leftarrow \{1\}$.
$\mathcal{I} \leftarrow \{2, 3, \ldots, K\}$.
for $(t = 2; t \leftarrow t + 1; t \leq r)$ **do**
  Define, $\tilde{\mathcal{C}}_i = \mathcal{C} \bigcup \{i\}, \forall i \in \mathcal{I}$.
  Run Algorithm 6.2 on matrices $\boldsymbol{L}_{\tilde{\mathcal{C}}_i, \tilde{\mathcal{C}}_i}$, $i \in \mathcal{I}$, with given $\delta \leftarrow \frac{\delta}{2r}$ to get $i_t^\star$.
  $\mathcal{C} \leftarrow \mathcal{C} \bigcup \{i_t^\star\}; \mathcal{I} \leftarrow \mathcal{I} \setminus \{i_t^\star\}$.
end for
Obtain estimators $\widehat{\boldsymbol{C}}, \widehat{\boldsymbol{W}}$ of $\boldsymbol{C}, \boldsymbol{W}$ by repeatedly sampling and averaging entries. Calculate the Nystrom extension $\widehat{\boldsymbol{L}} = \widehat{\boldsymbol{C}} \widehat{\boldsymbol{W}}^{-1} \widehat{\boldsymbol{C}}^\top$ (see lines 221 - 227).

Algorithm 6.2 Successive elimination on principal submatrices

Input: Square matrices $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_m$ of size $p \times p$, which share the same $p - 1 \times p - 1$ left principal submatrix; a failure probability $\delta > 0$; and a stochastic oracle $\mathcal{O}$

Output: An index

Set $t = 1$, and $\mathcal{S} = \{1, 2, \ldots, m\}$.
Sample each entry of the input matrices once.
while $|\mathcal{S}| > 1$ **do**
  Set $\delta_t = \frac{6\delta}{\pi^2 m t^2}$
  Let $\widehat{\sigma}^{\max} = \max_{k \in \mathcal{S}} \sigma_{\min}(\widehat{\boldsymbol{A}}_k)$ and let $k_\star$ be the index that attains argmax.
  For each $k \in \mathcal{S}$, define $\alpha_{t,k} = \frac{2 \log(2p/\delta_t)}{3 \min_{i,j} n_{i,j}(\widehat{\boldsymbol{A}}_k)} + \sqrt{\frac{\log(2p/\delta_t)}{2} \sum_{i,j} \frac{1}{n_{i,j}(\widehat{\boldsymbol{A}}_k)}}$
  For each index $k \in \mathcal{S}$, if $\widehat{\sigma}^{\max} - \widehat{\sigma}_{\min}(\widehat{\boldsymbol{A}}_k) \geq \alpha_{t,k^\star} + \alpha_{t,k}$ then do $\mathcal{S} \leftarrow \mathcal{S} \setminus \{k\}$.
  $t \leftarrow t + 1$
  Sample each entry of the matrices indexed by the indices in $\mathcal{S}$ once.
end while
Output k, where $k \in \mathcal{S}$.

6.1 Sample complexity of the R-PLANS algorithm

As can be seen from the R-PLANS algorithm, samples are consumed both in the successive elimination steps as well as during the construction of the Nyström extension. We shall now analyze both these steps next.

Sample complexity analysis of successive elimination. Before we provide a sample complexity of the analysis, we need a bound on the spectral norm of random matrices with 0 mean where each element is sampled possibly different number of times. This bound plays a key role in correctness of the successive elimination algorithm. The proof of this bound follows from matrix Bernstein inequality. We relegate the proof to the appendix due to lack of space.

Lemma 6.1. *Let $\hat{\boldsymbol{P}}$ be a $p \times p$ random matrix that is constructed as follows. For each index (i, j) independent of other indices, set $\hat{\boldsymbol{P}}_{i,j} = \frac{H_{i,j}}{n_{i,j}}$, where $H_{i,j}$ is a random variable drawn from the distribution $\mathrm{Binomial}(n_{i,j}, p_{i,j})$. Then, $||\hat{\boldsymbol{P}} - \boldsymbol{P}||_2 \leq \frac{2\log(2p/\delta)}{3\min_{i,j} n_{i,j}} + \sqrt{\frac{\log(2p/\delta)}{2}\sum_{i,j}\frac{1}{n_{i,j}}}$. Furthermore, if we denote by Δ the R.H.S. in the above bound, then $|\sigma_{\min}(\hat{\boldsymbol{P}}) - \sigma_{\min}(\boldsymbol{P})| \leq \Delta$.*

Lemma 6.2. *The successive elimination algorithm shown in Figure* (6.2) *on m square matrices of size $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_m$ each of size $p \times p$ outputs an index $i_\star$ such that, with probability at least $1 - \delta$, the matrix $\boldsymbol{A}_{i_\star}$ has the largest smallest singular value among all the input matrices. Let, $\Delta_{k,p} := \max_{j=1,\ldots,m} \sigma_{\min}(\boldsymbol{A}_j) - \sigma_{\min}(\boldsymbol{A}_k)$. Then number of queries to the stochastic oracle are*

$$\sum_{k=2}^{m} O\left(p^3 \log(2p\pi^2 m^2/3\Delta_{k,p}^2\delta)/\Delta_{k,p}^2\right) + O\left(p^4 \max_k \log(2p\pi^2 m^2/3\Delta_{k,p}^2\delta)/\Delta_{k,p}^2\right) \tag{11}$$

Sample complexity analysis of Nystrom extension. The following theorem tells us how many calls to a stochastic oracle are needed in order to guarantee that the Nystrom extension obtained by using matrices $\widehat{\boldsymbol{C}}, \widehat{\boldsymbol{W}}$ is accurate with high probability. The proof has been relegated to the appendix.

Theorem 6.3. *Consider the matrix $\widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}\widehat{\boldsymbol{C}}^\top$ which is the Nystrom extension constructed in step 10 of the R-PLANS algorithm. Given any $\delta \in (0, 1)$, with probability at least $1 - \delta$, $\left\|\boldsymbol{C}\boldsymbol{W}^{-1}\boldsymbol{C}^\top - \widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}\widehat{\boldsymbol{C}}^\top\right\|_{\max} \leq \epsilon$ after making a total of $Krm_1 + r^2 m_2$ number of oracle calls to a stochastic oracle, where m_1, m_2 are given in equations* (8)*.*

The following corollary follows directly from theorem (6.3), and lemma (6.2).

Corollary 6.4. *The R-PLANS algorithm outputs an (ϵ, δ) good arm after making at most*

$$\max\left(Krm_1 + r^2 m_2, \sum_{p=1}^{r}\sum_{k=2}^{K-r} \tilde{O}\left(\frac{p^3}{\Delta_{k,p}^2}\right) + \tilde{O}\left(p^4 \max_k(\frac{1}{\Delta_{k,p}^2})\right)\right)$$

number of calls to a stochastic oracle, where $\tilde{O}$ hides factors that are logarithmic in $K, r, \frac{1}{\delta}, 1/\Delta_{k,p}$.

7 Experiments

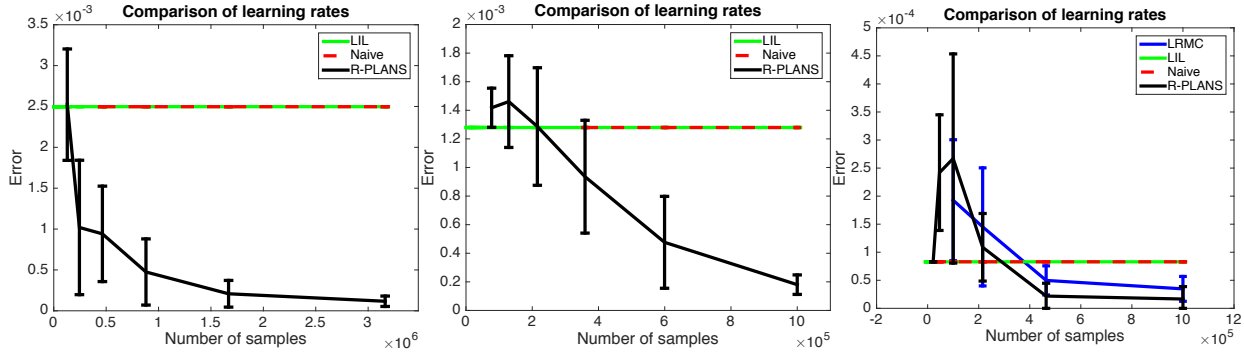

(a) ML-100K; $K = 800, r = 2$ (b) ML-100K; $K = 800, r = 4$ (c) ML-1M; $K = 200, r = 2$

Figure 1: Error of various algorithms with increasing budget. The error is defined as $\boldsymbol{L}_{\hat{i},\hat{j}} - \boldsymbol{L}_{i_\star,j_\star}$ where $(\hat{i}, \hat{j})$ is a pair of optimal choices as estimated by each algorithm. Note that the Naive and LIL' UCB have similar performances and do not improve with budget and both are outperformed by R-PLANS. This is because both Naive and LIL' UCB have few samples that they can use on each pair of movies. All experiments were repeated 10 times.

In the experiments demonstrated in this section we provide the different algorithms with increasing budget and measure the error of each algorithm in finding a pair of choices. The algorithms that we use for comparison are Naive, LiL'UCB [23] and LRMC using OptSpace [24]. The naive algorithm uniformly distributes equally the given budget to all $K(K+1)/2$ pairs of choices. LIL applies the LiL' UCB algorithm treating each pair of choices as an arm in a stochastic multi-armed bandit game.

In Figure (1) we demonstrate results on two Movie Lens datasets [25], namely ML-100K, ML-1M. These datasets consist of incomplete user-movie ratings. The task here is to recommend a pair of movies by using this user rating data. The datasets used in our experiments are generated by using the raw dataset and performing LRMC to get missing ratings, followed by thresholding the ratings to obtain binary values. The users were clustered as per their gender in Figures (1a), (1b) and as per their occupation in Figure (1b). We also sub-sampled the movies in our experiments. As can be seen from the figures the Naive and LIL'UCB algorithms have similar performance on all the datasets. On the ML-100K datasets LIL'UCB quickly finds a good pair of movies but fails to improve with an increase in the budget. To see why, observe that there are about 32×10^4 pairs of movies. The maximum budget here is on the order of 10^6. Therefore, Naive samples each of those pairs on an average at most four times. Since most entries in the matrix are of the order of 10^{-4}, Naive algorithm mostly sees 0s when sampling. The same thing happens with the LIL'UCB algorithm too; very few samples are available for each pair to improve its confidence bounds. This explains why the Naive and LIL' UCB algorithm have such poor and similar performances. In contrast, R-PLANS focusses most of the budget on a few select pairs and infers the value of other pairs via Nystrom extensio. This is why, in our experiments we see that R-PLANS finds good pair of movies quickly and finds even better pairs with increasing budget, outperforming all the other algorithms. R-PLANS is also better than LRMC, because we specifically exploit the SPSD structure in our matrix $\boldsymbol{L}$, which enables us to do better. We would like to mention that on ML-100K dataset the performance of LRMC was much inferior. This result and results on synthetic datasets can be found in the appendix

8 Conclusions

In this paper we introduced models for the preference learning problem and showed how to design active learning algorithms via reduction to completion of SPSD matrices. An interesting future work is to design algorithms and models when more than two choices are made each time. One can expect that techniques built in this paper adapted to tensor structure will be useful in such problems.

A Preliminaries

We shall repeat a proposition that was stated in the main paper for the sake of completeness.

Proposition A.1. *Let $\boldsymbol{L}$ be any SPSD matrix of size K. Given a subset $\mathcal{C} \subset \{1, 2, \ldots, K\}$, the columns of the matrix $\boldsymbol{L}$ indexed by the set $\mathcal{C}$ are independent iff the principal submatrix $\boldsymbol{L}_{\mathcal{C},\mathcal{C}}$ is non-degenerate, equivalently iff, $\lambda_{\min}(\boldsymbol{L}_{\mathcal{C},\mathcal{C}}) > 0$.*

We would also need the classical matrix Bernstein inequality, which we borrow from the work of Joel Tropp [26].

Theorem A.2. *Let $\boldsymbol{S}_1, \ldots, \boldsymbol{S}_n$ be independent, centered random matrices with dimension $d_1 \times d_2$ and assume that each one is uniformly bounded*

$$\mathbb{E}\boldsymbol{S}_k = 0, \|\boldsymbol{S}_k\| \leq L \text{for each } k = 1, \ldots, n.$$

Introduce the sum $\boldsymbol{Z} = \sum_{k=1}^{n} \boldsymbol{S}_k$, and let $\nu(\boldsymbol{Z})$ denote the matrix variance statistic of the sum:

$$\nu(\boldsymbol{Z}) = \max\left\{\left\|\mathbb{E}\boldsymbol{Z}\boldsymbol{Z}^\top\right\|, \left\|\mathbb{E}\boldsymbol{Z}^\top\boldsymbol{Z}\right\|\right\} \tag{12}$$

$$= \max\left\{\left\|\sum_{k=1}^{n}\mathbb{E}\boldsymbol{S}_k\boldsymbol{S}_k^\top\right\|, \left\|\sum_{k=1}^{n}\mathbb{E}\boldsymbol{S}_k^\top\boldsymbol{S}_k\right\|\right\} \tag{13}$$

Then,

$$\mathbb{P}(\|\boldsymbol{Z}\| \geq t) \leq (d_1 + d_2)\exp\left(-\frac{t^2/2}{\nu(\boldsymbol{Z}) + \frac{Lt}{3}}\right)$$

B Sample complexity of PLANS algorithm: Proof of Theorem 5.2

Theorem B.1. *If $\boldsymbol{L} \in \mathbb{R}^{K\times K}$ is an SPSD matrix of rank r, then the matrix $\hat{\boldsymbol{L}}$ output by the PLANS algorithm satisfies $\hat{\boldsymbol{L}} = \boldsymbol{L}$. Moreover, the number of oracle calls made by PLANS is at most $K(r+1)$. The sampling algorithm requires: $K + (K-1) + (K-2) + \ldots + (K-(r-1)) + (K-r) \leq (r+1)K$ samples from the matrix $\boldsymbol{L}$.*

Proof. PLANS checks one column at a time starting from the second column, and uses the test in line 5 to determine if the current column is independent of the previous columns. The validity of this test is guaranteed by proposition (A.1). Each such test needs just one additional sample corresponding to the index (c, c). If a column c is found to be independent of the columns $1, 2, \ldots, c-1$ then rest of the entries in column c are queried. Notice, that by now we have already queried all the columns and rows of matrix $\boldsymbol{L}$ indexed by the set $\mathcal{C}$, and also queried the element (c, c) in line 4. Hence we need to query only $K - |\mathcal{C}| - 1$ more entries in column c in order to have all the entries of column c. Combined with the fact that we query only r columns completely and in the worst case all the diagonal entries might be queried, we get the total query complexity to be $(K-1) + (K-2) + \ldots (K-r) + K \leq K(r+1)$. □

C Proof of Lemma 6.1

We begin by stating the lemma.

Lemma. *Let $\hat{\boldsymbol{P}}$ be a $p \times p$ random matrix that is constructed as follows. For each index (i, j) independent of other indices, set $\hat{\boldsymbol{P}}_{i,j} = \frac{H_{i,j}}{n_{i,j}}$, where $H_{i,j}$ is a random variable drawn from the distribution $\mathrm{Binomial}(n_{i,j}, p_{i,j})$. Let $Z = \hat{\boldsymbol{P}} - \boldsymbol{P}$. Then,*

$$||Z||_2 \leq \frac{2\log(2p/\delta)}{3\min_{i,j} n_{i,j}} + \sqrt{\frac{\log(2p/\delta)}{2}\sum_{i,j}\frac{1}{n_{i,j}}}. \tag{14}$$

Furthermore, if we denote by Δ the R.H.S. in Equation (14)*, then $|\sigma_{\min}(\hat{\boldsymbol{P}}) - \sigma_{\min}(\boldsymbol{P})| \leq \Delta$.*

Proof. Define, $\boldsymbol{S}_{i,j}^t = \frac{1}{n_{i,j}}(X_{i,j}^t - p_{i,j})\boldsymbol{E}_{i,j}$, where $\boldsymbol{E}_{i,j}$ is a $p \times p$ matrix with a 1 in the $(i,j)^{\text{th}}$ entry and 0 everywhere else, and $X_{i,j}^t$ is a random variable sampled from the distribution $\text{Bern}(p_{i,j})$. If $X_{i,j}^t$ are independent for all t, i, j, then it is easy to see that $Z = \sum_{i,j} \frac{1}{n_{i,j}} \sum_{t=1}^{n_{i,j}} \boldsymbol{S}_{i,j}^t$. Hence $\boldsymbol{S}$ is a sum of independent random matrices and this allows to apply matrix Bernstein type inequalities. In order to apply the matrix Bernstein inequality, we would need upper bound on maximum spectral norm of the summands, and an upper bound on the variance of $\boldsymbol{Z}$. We next bound these two quantities as follows,

$$||\boldsymbol{S}_{i,j}^t||_2 = ||\frac{1}{n_{i,j}}(X_{i,j}^t - p_{i,j})\boldsymbol{E}_{i,j}||_2 = \frac{1}{n_{i,j}}|X_{i,j}^t - p_{i,j}| \leq \frac{1}{n_{i,j}}. \tag{15}$$

To bound the variance of $\boldsymbol{Z}$ we proceed as follows

$$\nu(Z) = ||\sum_{i,j}\sum_{t=1}^{n_{i,j}} \mathbb{E}(\boldsymbol{S}_{i,j}^t)^\top \boldsymbol{S}_{i,j}^t|| \bigwedge ||\sum_{i,j}\sum_{t=1}^{n_{i,j}} \mathbb{E}\boldsymbol{S}_{i,j}^t(\boldsymbol{S}_{i,j}^t)^\top|| \tag{16}$$

Via elementary algebra and using the fact that $\text{Var}(X_{i,j}^t) = p_{i,j}(1-p_{i,j})$It is easy to see that,

$$\mathbb{E}(\boldsymbol{S}_{i,j}^t)^\top \boldsymbol{S}_{i,j}^t = \frac{1}{n_{i,j}^2}\mathbb{E}(X_{i,j}^t - p_{i,j})^2(\boldsymbol{E}_{i,j}^t)^\top \boldsymbol{E}_{i,j}^t \tag{17}$$

$$= \frac{1}{4n_{i,j}^2}\boldsymbol{E}_{i,i}. \tag{18}$$

Using similar calculations we get $\mathbb{E}\boldsymbol{S}_{i,j}^t(\boldsymbol{S}_{i,j}^t)^\top = \frac{1}{4n_{i,j}^2}\boldsymbol{E}_{j,j}$. Hence, $\nu(Z) = \sum_{i,j}\sum_{t=1}^{n_{i,j}} \frac{1}{4n_{i,j}^2} = \sum_{i,j}\frac{1}{4n_{i,j}}$. Applying matrix Bernstein, we get with probability at least $1-\delta$

$$||Z||_2 \leq \frac{2\log(2p/\delta)}{3\min_{i,j} n_{i,j}} + \sqrt{\frac{\log(2p/\delta)}{2}\sum_{i,j}\frac{1}{n_{i,j}}}. \tag{19}$$

The second part of the result follows immediately from Weyl's inequality which says that $|\sigma_{\min}(\hat{\boldsymbol{P}}) - \sigma_{\min}(\boldsymbol{P})| \leq ||\hat{\boldsymbol{P}} - \boldsymbol{P}|| = ||Z||$. □

D Sample complexity of successive elimination algorithm: Proof of Lemma 6.2

Lemma. *The successive elimination algorithm shown in Figure (6.2) on m square matrices of size $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_m$ each of size $p \times p$ outputs an index $i_\star$ such that, with probability at least $1-\delta$, the matrix $\boldsymbol{A}_{i_\star}$ has the largest smallest singular value among all the input matrices. The total number of queries to the stochastic oracle are*

$$\sum_{k=2}^{m} O\left(\frac{p^3\log(2p\pi^2m^2/3\Delta_k^2\delta)}{\Delta_k^2}\right) + O\left(p^4\max_k\left(\frac{\log(2p\pi^2m^2/3\Delta_k^2\delta)}{\Delta_k^2}\right)\right) \tag{20}$$

where $\Delta_{k,p} := \max_{j=1,\ldots,m}\sigma_{\min}(\boldsymbol{A}_j) - \sigma_{\min}(\boldsymbol{A}_k)$

Proof. Suppose matrix $\boldsymbol{A}_1$ has the largest smallest singular value. From lemma (C), we know that with probability at least $1-\delta_t$, $|\sigma_{\min}(\widehat{\boldsymbol{A}}_k) - \sigma_{\min}(\boldsymbol{A}_k)| \leq \frac{2\log(2p/\delta_t)}{3\min_{i,j} n_{i,j}(\boldsymbol{A})} + \sqrt{\frac{\log(2p/\delta_t)}{2}\sum_{i,j}\frac{1}{n_{i,j}(\boldsymbol{A})}}$. Hence, by union bound the probability that the matrix $\boldsymbol{A}_1$ is eliminated in one of the rounds is at most $\sum_t\sum_{k=1}^m \delta_t \leq \sum_{t=1}^{\max}\sum_{k=1}^m \frac{6\delta}{\pi^2mt^2} = \delta$. This proves that the successive elimination step identifies the matrix with the largest smallest singular value.

An arm k is eliminated in round t if $\alpha_{t,1} + \alpha_{t,k} \leq \hat{\sigma}_t^{\max} - \sigma_{\min}(\widehat{\boldsymbol{A}}_k)$. By definition,

$$\Delta_{k,p} - (\alpha_{t,1} + \alpha_{t,k}) = (\sigma_{\min}(\boldsymbol{A}_1) - \alpha_{t,1}) - (\sigma_{\min}(\boldsymbol{A}_k) + \alpha_{t,k}) \geq \sigma_{\min}(\hat{\boldsymbol{A}}_1) - \sigma_{\min}(\boldsymbol{A}_k) \geq \alpha_{t,1} + \alpha_{t,k} \tag{21}$$

That is if $\alpha_{t,1}+\alpha_{t,k} \leq \frac{\Delta_{k,p}}{2}$, then arm k is eliminated in round t. By construction, since in round t each element in each of the surviving set of matrices has been queried at least t times, we can say that $\alpha_{t,j} \leq \frac{2\log(2p/\delta_t)}{3t} + \sqrt{\frac{p^2\log(2p/\delta_t)}{2t}}$ for any index j corresponding to the set of surviving arms. Hence arm k gets eliminated after

$$t_k = O\left(\frac{p^2\log(2p\pi^2m^2/3\Delta_{k,p}^2\delta)}{\Delta_{k,p}^2}\right) \tag{22}$$

In each round t the number of queries made are $O(p)$ for each of the m matrices corresponding to the row and column which is different among them, and $O(p^2)$ corresponding to the left $p-1 \times p-1$ submatrix that is common to all of the matrices $A_1, \ldots, A_m$. Hence, the total number of queries to the stochastic oracle is

$$p\sum_{k=2}^{m} t_k + p^2 \max_k t_k = \sum_{k=2}^{m} O\left(\frac{p^3\log(2p\pi^2m^2/3\Delta_{k,p}^2\delta)}{\Delta_{k,p}^2}\right) + O\left(p^4 \max_k\left(\frac{\log(2p\pi^2m^2/3\Delta_{k,p}^2\delta)}{\Delta_{k,p}^2}\right)\right) \qquad \square$$

E Proof of Nystrom method

In this supplementary material we provide a proof of Nystrom extension in $\max$ norm when we use a stochastic oracle to obtain estimators $\widehat{\boldsymbol{C}}, \widehat{\boldsymbol{W}}$ of the matrices $\boldsymbol{C}, \boldsymbol{W}$. The question that we are interested in is how good is the estimate of the Nystrom extension obtained using matrices $\widehat{\boldsymbol{C}}, \widehat{\boldsymbol{W}}$ w.r.t. the Nystrom extension obtained using matrices $\boldsymbol{C}, \boldsymbol{W}$. This is answered in the theorem below.

Theorem E.1. *Suppose the matrix $\boldsymbol{W}$ is an invertible $r \times r$ matrix. Suppose, by multiple calls to a stochastic oracle we construct estimators $\widehat{\boldsymbol{C}}, \widehat{\boldsymbol{W}}$ of $\boldsymbol{C}, \boldsymbol{W}$. Now, consider the matrix $\widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}\widehat{\boldsymbol{C}}^\top$ as an estimate $\boldsymbol{C}\boldsymbol{W}^{-1}\boldsymbol{C}^\top$. Given any $\delta \in (0,1)$, with probability atleast $1-\delta$,*

$$\left\|\boldsymbol{C}\boldsymbol{W}^{-1}\boldsymbol{C}^\top - \widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}\widehat{\boldsymbol{C}}^\top\right\|_{\max} \leq \epsilon$$

after making M number of oracle calls to a stochastic oracle, where

$$M \geq 100C_1(W,C)\log(2Kr/\delta)\max\left(\frac{Kr^{7/2}}{\epsilon}, \frac{Kr^3}{\epsilon^2}\right) + 200C_2(W,C)\log(2r/\delta)\max\left(\frac{r^5}{\epsilon}, \frac{r^7}{\epsilon^2}\right)$$

where $C_1(\boldsymbol{W},\boldsymbol{C})$ and $C_2(\boldsymbol{W},\boldsymbol{C})$ are given by the following equations

$$C_1(\boldsymbol{W},\boldsymbol{C}) = \max\left(\left\|\boldsymbol{W}^{-1}\boldsymbol{C}^\top\right\|_{\max}, \left\|\boldsymbol{W}^{-1}\boldsymbol{C}^\top\right\|_{\max}^2, \left\|\boldsymbol{W}^{-1}\right\|_{\max}, \left\|\boldsymbol{C}\boldsymbol{W}^{-1}\right\|_1^2, \left\|\boldsymbol{W}^{-1}\right\|_2\left\|\boldsymbol{W}^{-1}\right\|_{\max}\right)$$
$$C_2(\boldsymbol{W},\boldsymbol{C}) = \max\left(\left\|\boldsymbol{W}^{-1}\right\|_2^2\left\|\boldsymbol{W}^{-1}\right\|_{\max}^2, \left\|\boldsymbol{W}^{-1}\right\|_2\left\|\boldsymbol{W}^{-1}\right\|_{\max}, \left\|\boldsymbol{W}^{-1}\right\|_2, \left\|\boldsymbol{W}^{-1}\right\|_2^2\right)$$

Our proof proceeds by a series of lemmas, which we state next.

Lemma E.2.

$$\left\|\boldsymbol{C}\boldsymbol{W}^{-1}\boldsymbol{C}^\top - \widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}\widehat{\boldsymbol{C}}^\top\right\|_{\max} \leq \left\|(\boldsymbol{C}-\widehat{\boldsymbol{C}})\boldsymbol{W}^{-1}\boldsymbol{C}^\top\right\|_{\max} + \left\|\widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}(\boldsymbol{C}-\widehat{\boldsymbol{C}})^\top\right\|_{\max} + \left\|\widehat{\boldsymbol{C}}(\boldsymbol{W}^{-1}-\widehat{\boldsymbol{W}}^{-1})\boldsymbol{C}^\top\right\|_{\max}$$

Proof.

$$\begin{aligned}
\left\|\boldsymbol{C}\boldsymbol{W}^{-1}\boldsymbol{C}^\top - \widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}\widehat{\boldsymbol{C}}^\top\right\|_{\max} &= \left\|\boldsymbol{C}\boldsymbol{W}^{-1}\boldsymbol{C}^\top - \widehat{\boldsymbol{C}}\boldsymbol{W}^{-1}\boldsymbol{C}^\top + \widehat{\boldsymbol{C}}\boldsymbol{W}^{-1}\boldsymbol{C}^\top - \widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}\widehat{\boldsymbol{C}}^\top\right\|_{\max} \\
&\leq \left\|\boldsymbol{C}\boldsymbol{W}^{-1}\boldsymbol{C}^\top - \widehat{\boldsymbol{C}}\boldsymbol{W}^{-1}\boldsymbol{C}^\top\right\|_{\max} + \left\|\widehat{\boldsymbol{C}}\boldsymbol{W}^{-1}\boldsymbol{C}^\top - \widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}\widehat{\boldsymbol{C}}^\top\right\|_{\max} \\
&= \left\|\boldsymbol{C}\boldsymbol{W}^{-1}\boldsymbol{C}^\top - \widehat{\boldsymbol{C}}\boldsymbol{W}^{-1}\boldsymbol{C}^\top\right\|_{\max} + \\
&\quad \left\|\widehat{\boldsymbol{C}}\boldsymbol{W}^{-1}\boldsymbol{C}^\top - \widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}\boldsymbol{C}^\top + \widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}\boldsymbol{C}^\top - \widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}\widehat{\boldsymbol{C}}^\top\right\|_{\max} \\
&\leq \left\|\boldsymbol{C}\boldsymbol{W}^{-1}\boldsymbol{C}^\top - \widehat{\boldsymbol{C}}\boldsymbol{W}^{-1}\boldsymbol{C}^\top\right\|_{\max} + \left\|\widehat{\boldsymbol{C}}W^{-1}\boldsymbol{C}^\top - \widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}\boldsymbol{C}^\top\right\|_{\max} + \\
&\quad \left\|\widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}\boldsymbol{C}^\top - \widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}\widehat{\boldsymbol{C}}^\top\right\|_{\max} \\
&= \left\|(\boldsymbol{C} - \widehat{\boldsymbol{C}})\boldsymbol{W}^{-1}\boldsymbol{C}^\top\right\|_{\max} + \left\|\widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}(\boldsymbol{C} - \widehat{\boldsymbol{C}})^\top\right\|_{\max} + \left\|\widehat{\boldsymbol{C}}(\boldsymbol{W}^{-1} - \widehat{\boldsymbol{W}}^{-1})\boldsymbol{C}^\top\right\|_{\max}
\end{aligned}$$

□

In the following lemmas we shall bound the three terms that appear in the R.H.S of the bound of Lemma (E.2).

Lemma E.3.

$$\left\|(C - \hat{C})W^{-1}C^\top\right\|_{\max} \leq \frac{2||\boldsymbol{W}^{-1}\boldsymbol{C}^\top||_{\max}}{3m}\log(2Kr/\delta) + \sqrt{\frac{r\,||\boldsymbol{W}^{-1}\boldsymbol{C}^\top||^2_{\max}\log(2Kr/\delta)}{2m}} \tag{23}$$

Proof. Let $\boldsymbol{M} = \boldsymbol{W}^{-1}\boldsymbol{C}^\top$, then $\left\|(\boldsymbol{C} - \widehat{\boldsymbol{C}})\boldsymbol{W}^{-1}\boldsymbol{C}^\top\right\|_{\max} = \left\|(\boldsymbol{C} - \widehat{\boldsymbol{C}})\boldsymbol{M}\right\|_{\max}$. By the definition of max norm we have

$$\left\|(\boldsymbol{C} - \widehat{\boldsymbol{C}})M\right\|_{\max} = \max_{i,j}\left|\sum_{p=1}^{l}(\boldsymbol{C} - \widehat{\boldsymbol{C}})_{i,p}\boldsymbol{M}_{p,j}\right|$$

Fix a pair of indices (i, j), and consider the expression $\left|\sum_{p=1}^{l}(\boldsymbol{C} - \widehat{\boldsymbol{C}})_{i,p}\boldsymbol{M}_{p,j}\right|$

Define $r_{i,p} = (\boldsymbol{C} - \widehat{\boldsymbol{C}})_{i,p}$. By definition of $r_{i,p}$ we can write $r_{i,p} = \frac{1}{m}\sum_{t=1}^{m} r^t_{i,p}$, where $r^t_{i,p}$ are a set of independent random variables with mean 0 and variance at most $1/4$. This decomposition combined with scalar Bernstein inequality gives that with probability at least $1 - \delta$

$$\begin{aligned}
\left|\sum_{p=1}^{l}(\widehat{\boldsymbol{C}} - \widehat{\boldsymbol{C}})_{i,p}\boldsymbol{M}_{p,j}\right| &= \left|\sum_{p=1}^{l} r_{i,p}\boldsymbol{M}_{p,j}\right| \\
&= \left|\sum_{p=1}^{l}\sum_{t=1}^{m}\frac{1}{m}r^t_{i,p}\boldsymbol{M}_{p,j}\right| \\
&\leq \frac{2||M||_{\max}}{3m}\log(2/\delta) + \sqrt{\frac{r\,||M||^2_{\max}\log(2/\delta)}{2m}}
\end{aligned}$$

Applying a union bound over all possible Kr choices of index pairs (i, j), we get the desired result. □

Before we establish bounds on the remaining two terms in the RHS of Lemma (E.2) we state and prove a simple proposition that will be used at many places in the rest of the proof.

Proposition E.4. *For any two real matrices $M_1 \in \mathbb{R}^{n_1\times n_2}, M_2 \in \mathbb{R}^{n_2\times n_3}$ the following set of inequalities are true:*

1. $\|\boldsymbol{M}_1\boldsymbol{M}_2\|_{\max} \leq \|\boldsymbol{M}_1\|_{\max}\|\boldsymbol{M}_2\|_1$

2. $\|\boldsymbol{M}_1\boldsymbol{M}_2\|_{\max} \leq \left\|\boldsymbol{M}_1^\top\right\|_1 \|\boldsymbol{M}_2\|_{\max}$
3. $\|\boldsymbol{M}_1\boldsymbol{M}_2\|_{\max} \leq \|\boldsymbol{M}_1\|_2 \|\boldsymbol{M}_2\|_{\max}$
4. $\|\boldsymbol{M}_1\boldsymbol{M}_2\|_{\max} \leq \|\boldsymbol{M}_2\|_2 \|\boldsymbol{M}_1\|_{\max}$

where, the $\|\cdot\|_p$ is the induced p norm.

Proof. Let $\boldsymbol{e}_i$ denote the i^{th} canonical basis vectors in $\mathbb{R}^K$. We have,

$$\begin{aligned}
\|\boldsymbol{M}_1\boldsymbol{M}_2\|_{\max} &= \max_{i,j} \left|e_i^\top \boldsymbol{M}_1\boldsymbol{M}_2 e_j\right| \\
&\leq \max_{i,j} \left\|e_i^\top \boldsymbol{M}_1\right\|_{\max} \|\boldsymbol{M}_2 e_j\|_1 \\
&= \max_i \left\|e_i^\top \boldsymbol{M}_1\right\|_{\max} \max_i \|\boldsymbol{M}_2 e_j\|_1 \\
&= \|\boldsymbol{M}_1\|_{\max} \|\boldsymbol{M}_2\|_1 .
\end{aligned}$$

To obtain the first inequality above we used Holder's inequality and the last equality follows from the definition of $||\cdot||_1$ norm. To get the second inequality, we use the observations that $\|\boldsymbol{M}_1\boldsymbol{M}_2\|_{\max} = \left\|\boldsymbol{M}_2^\top \boldsymbol{M}_1^\top\right\|_{\max}$. Now applying the first inequality to this expression we get the desired result. Similar techniques yield the other two inequalities. □

Lemma E.5. *With probability at least $1-\delta$, we have*

$$\begin{aligned}
\left\|\widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}(\boldsymbol{C}-\widehat{\boldsymbol{C}})^\top\right\|_{\max} \leq \frac{r^2}{2m}\left(\left\|\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{W}^{-1}\right\|_{\max} + \|\boldsymbol{W}^{-1}\|_{\max}\right)\log(2Kr/\delta) + \\
r^2\left\|\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{W}^{-1}\right\|_{\max}\sqrt{\frac{\log(2Kr/\delta)}{2m}} + r\left\|\boldsymbol{C}\boldsymbol{W}^{-1}\right\|_1\sqrt{\frac{\log(2Kr/\delta)}{2m}}
\end{aligned}$$

Proof.

$$\begin{aligned}
\left\|\widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}(\boldsymbol{C}-\widehat{\boldsymbol{C}})^\top\right\|_{\max} &\leq \left\|(\widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{C}\boldsymbol{W}^{-1}+\boldsymbol{C}\boldsymbol{W}^{-1})(\boldsymbol{C}-\widehat{\boldsymbol{C}})^\top\right\|_{\max} \\
&\overset{(a)}{\leq} \left\|(\widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{C}\boldsymbol{W}^{-1})(\boldsymbol{C}-\widehat{\boldsymbol{C}})^\top\right\|_{\max} + \left\|\boldsymbol{C}\boldsymbol{W}^{-1}(\boldsymbol{C}-\widehat{\boldsymbol{C}})^\top\right\|_{\max} \\
&\overset{(b)}{\leq} \left\|\widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{C}\boldsymbol{W}^{-1}\right\|_{\max}\left\|(\boldsymbol{C}-\widehat{\boldsymbol{C}})^\top\right\|_1 + \left\|\boldsymbol{C}\boldsymbol{W}^{-1}\right\|_{\max}\left\|(\boldsymbol{C}-\widehat{\boldsymbol{C}})^\top\right\|_1 \qquad (24)
\end{aligned}$$

To obtain inequality (a) we used triangle inequality for matrix norms, and to obtain inequality (b) we used Proposition (E.4). We next upper bound the first term in the R.H.S. of Equation (24).

We bound the term $\left\|\widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{C}\boldsymbol{W}^{-1}\right\|_{\max}$ next.

$$\begin{aligned}
\left\|\widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{C}\boldsymbol{W}^{-1}\right\|_{\max} &\leq \left\|\widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{C}\widehat{\boldsymbol{W}}^{-1}+\boldsymbol{C}\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{C}\boldsymbol{W}^{-1}\right\|_{\max} \\
&\leq \left\|\widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{C}\widehat{\boldsymbol{W}}^{-1}\right\|_{\max} + \left\|\boldsymbol{C}\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{C}\boldsymbol{W}^{-1}\right\|_{\max} \\
&= \left\|(\widehat{\boldsymbol{C}}-\boldsymbol{C})\widehat{\boldsymbol{W}}^{-1}\right\|_{\max} + \left\|\boldsymbol{C}(\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{W}^{-1})\right\|_{\max} \\
&\overset{(a)}{\leq} \left\|(\widehat{\boldsymbol{C}}-\boldsymbol{C})^\top\right\|_1\left\|\widehat{\boldsymbol{W}}^{-1}\right\|_{\max} + \|\boldsymbol{C}^\top\|_1\left\|\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{W}^{-1}\right\|_{\max} \qquad (25)
\end{aligned}$$

We used Proposition (E.4) to obtain inequality (a). Combining Equations (24) and (25) we get,

$$\begin{aligned}
\left\|\widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}(\boldsymbol{C}-\widehat{\boldsymbol{C}})^\top\right\|_{\max} &\leq \left\|(\widehat{\boldsymbol{C}}-\boldsymbol{C})^\top\right\|_1\left(\left\|(\widehat{\boldsymbol{C}}-\boldsymbol{C})^\top\right\|_1\left\|\widehat{\boldsymbol{W}}^{-1}\right\|_{\max} + \|\boldsymbol{C}^\top\|_1\left\|\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{W}^{-1}\right\|_{\max} + \|\boldsymbol{C}\boldsymbol{W}^{-1}\|_{\max}\right) \\
&= \left\|(\widehat{\boldsymbol{C}}-\boldsymbol{C})^\top\right\|_1^2\left\|\widehat{\boldsymbol{W}}^{-1}\right\|_{\max} + \left\|(\widehat{\boldsymbol{C}}-\boldsymbol{C})^\top\right\|_1\|\boldsymbol{C}^\top\|_1\left\|\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{W}^{-1}\right\|_{\max} + \\
&\quad \left\|(\widehat{\boldsymbol{C}}-\boldsymbol{C})^\top\right\|_1\|\boldsymbol{C}\boldsymbol{W}^{-1}\|_{\max} \qquad (26)
\end{aligned}$$

Since all the entries of the matrix $\boldsymbol{C}$ are probabilities we have $\|\boldsymbol{C}\|_{\max} \leq 1$ and $\left\|\boldsymbol{C}^\top\right\|_1 \leq r$. Moreover, since each entry of the matrix $\hat{\boldsymbol{C}} - \boldsymbol{C}$ is the average of m independent random variables with mean 0, and each bounded between $[-1, 1]$, by Hoeffding's inequality and union bound, we get that with probability at least $1 - \delta$

$$\left\|(\widehat{\boldsymbol{C}} - \boldsymbol{C})^\top\right\|_1 \leq r\sqrt{\frac{\log(2Kr/\delta)}{2m}} \tag{27}$$

□

The next proposition takes the first steps towards obtaining an upper bound on $\left\|\hat{C}(W^{-1} - \hat{W}^{-1})C^\top\right\|_{\max}$

Proposition E.6.

$$\left\|\widehat{\boldsymbol{C}}(\boldsymbol{W}^{-1} - \widehat{\boldsymbol{W}}^{-1})\boldsymbol{C}^\top\right\|_{\max} \leq \min\left\{r^2\left\|\boldsymbol{W}^{-1} - \widehat{\boldsymbol{W}}^{-1}\right\|_{\max}, r\left\|\boldsymbol{W}^{-1} - \widehat{\boldsymbol{W}}^{-1}\right\|_1\right\}$$

Proof.

$$\begin{aligned}
\left\|\widehat{\boldsymbol{C}}(\boldsymbol{W}^{-1} - \widehat{\boldsymbol{W}}^{-1})\boldsymbol{C}^\top\right\|_{\max} &\overset{(a)}{\leq} \left\|\widehat{\boldsymbol{C}}(\boldsymbol{W}^{-1} - \widehat{\boldsymbol{W}}^{-1})\right\|_{\max}\left\|\boldsymbol{C}^\top\right\|_1 \\
&\overset{(b)}{\leq} r\left\|\widehat{\boldsymbol{C}}(\boldsymbol{W}^{-1} - \widehat{\boldsymbol{W}}^{-1})\right\|_{\max} \\
&\overset{(c)}{\leq} \min\left\{r^2\left\|\boldsymbol{W}^{-1} - \widehat{\boldsymbol{W}}^{-1}\right\|_{\max}, r\left\|\boldsymbol{W}^{-1} - \widehat{\boldsymbol{W}}^{-1}\right\|_1\right\}
\end{aligned} \tag{28}$$

In the above bunch of inequalities (a) and (c) we used Proposition (E.4) and to obtain inequality (b) we used the fact that $||\boldsymbol{C}^\top||_{\max} \leq r$. □

Hence, we need to bound $\left\|\boldsymbol{W}^{-1} - \widehat{\boldsymbol{W}}^{-1}\right\|_{\max}$ and $\left\|\boldsymbol{W}^{-1} - \widehat{\boldsymbol{W}}^{-1}\right\|_1$.

Let us define $\hat{\boldsymbol{W}} = \boldsymbol{W} + \boldsymbol{E}_W$ where $\boldsymbol{E}_W$ is the error-matrix and $\hat{\boldsymbol{W}}$ is the sample average of m independent samples of a random matrix where $\mathbb{E}\hat{\boldsymbol{W}}_k(i,j) = \boldsymbol{W}(i,j)$.

Lemma E.7. *Let us define $\hat{\boldsymbol{W}} - \boldsymbol{W} = \boldsymbol{E}_W$. Suppose, $\left\|\boldsymbol{W}^{-1}\boldsymbol{E}_W\right\|_2 \leq \frac{1}{2}$, then*

$$\left\|\hat{\boldsymbol{W}}^{-1} - \boldsymbol{W}^{-1}\right\|_{\max} \leq 2\left\|\boldsymbol{W}^{-1}\right\|_2\|\boldsymbol{E}_W\|_2\left\|\boldsymbol{W}^{-1}\right\|_{\max}$$

.

Proof. Since $\left\|\boldsymbol{W}^{-1}\boldsymbol{E}_W\right\|_2 < 1$, we can apply the Taylor series expansion:

$$(\boldsymbol{W} + \boldsymbol{E}_W)^{-1} = \boldsymbol{W}^{-1} - \boldsymbol{W}^{-1}\boldsymbol{E}_W\boldsymbol{W}^{-1} + \boldsymbol{W}^{-1}\boldsymbol{E}_W\boldsymbol{W}^{-1}\boldsymbol{E}_W\boldsymbol{W}^{-1} + \cdots$$

Therefore:

$$\begin{aligned}
\left\|\hat{\boldsymbol{W}}^{-1} - \boldsymbol{W}^{-1}\right\|_{\max} &= \left\|\boldsymbol{W}^{-1} - \boldsymbol{W}^{-1}\boldsymbol{E}_W\boldsymbol{W}^{-1} + \boldsymbol{W}^{-1}\boldsymbol{E}_W\boldsymbol{W}^{-1}\boldsymbol{E}_W\boldsymbol{W}^{-1} + \cdots - \boldsymbol{W}^{-1}\right\|_{\max} \\
&\overset{(a)}{\leq} \left\|\boldsymbol{W}^{-1}\boldsymbol{E}_W\boldsymbol{W}^{-1}\right\|_{\max} + \left\|\boldsymbol{W}^{-1}\boldsymbol{E}_W\boldsymbol{W}^{-1}\boldsymbol{E}_W\boldsymbol{W}^{-1}\right\|_{\max} + \cdots \\
&\overset{(b)}{\leq} \left\|\boldsymbol{W}^{-1}\boldsymbol{E}_W\right\|_2\left\|W^{-1}\right\|_{\max} + \left\|\boldsymbol{W}^{-1}\boldsymbol{E}_W\right\|_2^2\left\|W^{-1}\right\|_{\max} + \ldots \\
&\overset{(c)}{\leq} 2\left\|\boldsymbol{W}^{-1}\right\|_2\|\boldsymbol{E}_W\|_2\left\|\boldsymbol{W}^{-1}\right\|_{\max}
\end{aligned}$$

To obtain the last inequality we used the hypothesis of the lemma, and to obtain inequality (a) we used the triangle inequality for norms, and to obtain inequality (b) we used proposition (E.4). Inequlaity (c) follows from the triangle inequality. □

Thanks to Lemma (E.7) and proposition (E.6) we know that $\left\|\widehat{\boldsymbol{C}}(\boldsymbol{W}^{-1}-\widehat{\boldsymbol{W}}^{-1})\boldsymbol{C}^\top\right\|_{\max} \leq r^2\epsilon$. We now need to guarantee that the hypothesis of lemma (E.7) applies. The next lemma helps in doing that.

Lemma E.8. *With probability at least* $1-\delta$ *we have*

$$\|\boldsymbol{E}_W\| = \left\|\hat{\boldsymbol{W}}-\boldsymbol{W}\right\| \leq \frac{2r}{3m}\log(2r/\delta) + \sqrt{\frac{r\log(2r/\delta)}{2m}} \tag{29}$$

Proof. The proof is via matrix Bernstein inequality. By the definition of $\hat{\boldsymbol{W}}$, we know that $\hat{\boldsymbol{W}}-\boldsymbol{W} = \frac{1}{m}\sum(\boldsymbol{W}_i-\boldsymbol{W})$, where $\hat{\boldsymbol{W}}$ is $0-1$ random matrix where the $(i,j)^{\text{th}}$ entry of the matrix $\hat{\boldsymbol{W}}$ is a single Bernoulli sample sampled from $\mathrm{Bern}(\boldsymbol{W}_{i,j})$. For notational convenience denote $Z_i := \frac{1}{m}\hat{\boldsymbol{W}}_i - \boldsymbol{W}$ This makes $\hat{\boldsymbol{W}}-\boldsymbol{W} = \frac{1}{m}\sum\boldsymbol{W}_i - \boldsymbol{W}$ an average of m independent random matrices each of whose entry is a 0 mean random variable with variance at most $1/4$, with each entry being in $[-1,1]$. In order to apply the matrix Bernstein inequality we need to upper bound ν, L (see Theorem (A.2)), which we do next.

$$\left\|\frac{1}{m}(\hat{\boldsymbol{W}}_i-\boldsymbol{W})\right\|_2 \leq \frac{1}{m}\sqrt{r^2} = \frac{r}{m}. \tag{30}$$

In the above inequality we used the fact that each entry of $(\hat{\boldsymbol{W}}_i-\boldsymbol{W})$ is between $[-1,1]$ and hence the spectral norm of this matrix is at most $\sqrt{r^2}$. We next bound the parameter ν.

$$\nu = \frac{1}{m^2}\max\left\{\left\|\sum_i \mathbb{E}\boldsymbol{Z}_i\boldsymbol{Z}_i^\top\right\|, \left\|\sum_i \mathbb{E}\boldsymbol{Z}_i^\top\boldsymbol{Z}_i\right\|\right\} \tag{31}$$

It is not hard to see that the matrix $\mathbb{E}\boldsymbol{Z}_i\boldsymbol{Z}_i^\top$ is a diagonal matrix, where each diagonal entry is at most $\frac{l}{4}$. The same holds true for $\mathbb{E}\boldsymbol{Z}_i\boldsymbol{Z}_i^\top$. Putting this back in Equation (31) we get $\nu \leq \frac{r}{4m}$. Putting $L = \frac{r}{m}$ and $\nu = \frac{r}{4m}$, we get

$$\left\|\hat{\boldsymbol{W}}-\boldsymbol{W}\right\| \leq \frac{2r}{3m}\log(2r/\delta) + \sqrt{\frac{r\log(2r/\delta)}{2m}} \tag{32}$$

□

We are now ready to establish the following bound

Lemma E.9. *Assuming that* $m \geq m_0 := \frac{4r\|\boldsymbol{W}^{-1}\|}{3} + 2r\log(2r/\delta)\left\|\boldsymbol{W}^{-1}\right\|_2^2$*, with probability at least* $1-\delta$ *we will have*

$$\left\|\widehat{\boldsymbol{C}}(\boldsymbol{W}^{-1}-\widehat{\boldsymbol{W}}^{-1})\boldsymbol{C}^\top\right\|_{\max} \leq 2r^2\left\|\boldsymbol{W}^{-1}\right\|_2\left\|\boldsymbol{W}^{-1}\right\|_{\max}\left(\frac{2r}{3m}\log(2r/\delta) + \sqrt{\frac{r\log(2r/\delta)}{2m}}\right). \tag{33}$$

Proof.

$$\begin{aligned}
\left\|\widehat{\boldsymbol{C}}(\boldsymbol{W}^{-1}-\widehat{\boldsymbol{W}}^{-1})\boldsymbol{C}^\top\right\|_{\max} &\overset{\text{(a)}}{\leq} r^2\left\|\boldsymbol{W}^{-1}-\hat{\boldsymbol{W}}^{-1}\right\|_{\max} \\
&\overset{\text{(b)}}{\leq} 2r^2\left\|\boldsymbol{W}^{-1}\boldsymbol{E}_{\boldsymbol{W}}\right\|_2\left\|\boldsymbol{W}^{-1}\right\|_{\max} \\
&\overset{\text{(c)}}{\leq} 2r^2\left\|\boldsymbol{W}^{-1}\right\|_2\|\boldsymbol{E}_{\boldsymbol{W}}\|_2\left\|\boldsymbol{W}^{-1}\right\|_{\max} \\
&\overset{\text{(d)}}{\leq} 2r^2\left\|\boldsymbol{W}^{-1}\right\|_2\left\|\boldsymbol{W}^{-1}\right\|_{\max}\left(\frac{2r}{3m}\log(2r/\delta) + \sqrt{\frac{r\log(2r/\delta)}{2m}}\right)
\end{aligned}$$

□

To obtain inequality (a) above we used proposition (E.6), to obtain inequality (b) we used lemma (E.7), and finally to obtain inequality (c) we used the fact that matrix 2-norms are submultiplicative.

With this we now have bounds on all the necessary quantities. The proof of our theorem essentially requires us to put all these terms together.

F Proof of Theorem (E.1)

Since we need the total error in $\max$ norm to be at most ϵ, we will enforce that each term of our expression be atmost $\frac{\epsilon}{10}$. From lemma (E.2) we know that the maxnorm is the sum of three terms. Let us call the three terms in the R.H.S. of Lemma (E.2) T_1, T_2, T_3 repsectively. We then have that if we have m_1 number of copies of the matrix $\boldsymbol{C}$, where

$$m_1 \geq \frac{20 \left\|\boldsymbol{W}^{-1}\boldsymbol{C}^\top\right\|_{\max} \log(2Kr/\delta)}{3\epsilon} \bigwedge \frac{100r \left\|\boldsymbol{W}^{-1}\boldsymbol{C}^\top\right\|_{\max}^2 \log(2Kr/\delta)}{2\epsilon^2} \tag{34}$$

then $T_1 \leq \epsilon/5$. Next we look at T_3. From lemma (E.9) it is easy to see that we need m_3 independent copies of the matrix $\boldsymbol{W}$ so that $T_3 \leq \epsilon/5$, where m_3 is equal to

$$m_3 \geq \frac{40r^3 \left\|\boldsymbol{W}^{-1}\right\|_2 \left\|\boldsymbol{W}^{-1}\right\|_{\max} \log(2r/\delta)}{3\epsilon} \bigwedge \frac{400r^5 \left\|\boldsymbol{W}^{-1}\right\|_2^2 \left\|\boldsymbol{W}^{-1}\right\|_{\max}^2 \log(2r/\delta)}{2\epsilon^2} \tag{35}$$

Finally we now look at T_2. Combining lemma (E.5), and lemma (E.7) and (E.8) and after some elementary algebraic calculations we get that we need m_2 independent copies of the matrix $\boldsymbol{C}$ and $\boldsymbol{W}$ to get $T_2 \leq \frac{3\epsilon}{5}$, where m_2 is

$$m_2 \geq 100 \max(\left\|\boldsymbol{W}^{-1}\right\|_{\max}, \left\|\boldsymbol{C}\boldsymbol{W}^{-1}\right\|_1^2, \left\|\boldsymbol{W}^{-1}\right\|_2 \left\|\boldsymbol{W}^{-1}\right\|_{\max}) \log(2Kr/\delta) \left(\frac{r^{5/2}}{\epsilon}, \frac{r^2}{\epsilon^2}\right) \tag{36}$$

The number of calls to stochastic oracle is $r^2(m_0 + m_3) + Kr(m_1 + m_2)$, where m_0 is the number as stated in Lemma (E.9). Using the above derived bounds for $m_0 + m_1, m_2, m_3$ we get

$$\begin{aligned} Kr(m_1 + m_2) + r^2(m_0 + m_3) \geq 100 \log(2Kr/\delta) C_1(W, C) \max\left(\frac{Kr^{7/2}}{\epsilon}, \frac{Kr^3}{\epsilon^2}\right) + \\ 200 C_2(W, C) \log(2r/\delta) \max\left(\frac{r^5}{\epsilon}, \frac{r^7}{\epsilon^2}\right) \end{aligned}$$

where $C_1(\boldsymbol{W}, \boldsymbol{C})$ and $C_2(\boldsymbol{W}, \boldsymbol{C})$ are given by the following equations

$$\begin{aligned} C_1(\boldsymbol{W}, \boldsymbol{C}) &= \max\left(\left\|\boldsymbol{W}^{-1}\boldsymbol{C}^\top\right\|_{\max}, \left\|\boldsymbol{W}^{-1}\boldsymbol{C}^\top\right\|_{\max}^2, \left\|\boldsymbol{W}^{-1}\right\|_{\max}, \left\|\boldsymbol{C}\boldsymbol{W}^{-1}\right\|_1^2, \left\|\boldsymbol{W}^{-1}\right\|_2 \left\|\boldsymbol{W}^{-1}\right\|_{\max}\right) \\ C_2(\boldsymbol{W}, \boldsymbol{C}) &= \max\left(\left\|\boldsymbol{W}^{-1}\right\|_2^2 \left\|\boldsymbol{W}^{-1}\right\|_{\max}^2, \left\|\boldsymbol{W}^{-1}\right\|_2 \left\|\boldsymbol{W}^{-1}\right\|_{\max}, \left\|\boldsymbol{W}^{-1}\right\|_2, \left\|\boldsymbol{W}^{-1}\right\|_2^2\right) \end{aligned}$$

G Additional experimental results: Comparison with LRMC on Movie Lens datasets

First we present the results on the synthetic dataset. To generate a low-rank matrix, we take a random matrix in $\boldsymbol{L}_1 = [0, 1]^{K \times r}$ and then define $\boldsymbol{L}_2 = \boldsymbol{L}_1 \boldsymbol{L}_1^\top$. Then get $\boldsymbol{L} = \boldsymbol{L}_2 / \max_{i,j}(\boldsymbol{L}_2)_{i,j}$. This matrix $\boldsymbol{L}$ will be $K \times K$ and have rank r.

In Figure 3, you can find the comparison of LRMC and R-PLANS on the ML-100K dataset.

References

[1] Moshe Babaioff, Yogeshwer Sharma, and Aleksandrs Slivkins. Characterizing truthful multi-armed bandit mechanisms. In *ACM-EC*. ACM, 2009.

[2] David Hensher, Jordan Louviere, and Joffre Swait. Combining sources of preference data. *Journal of Econometrics*, 89(1), 1998.

[3] Jordan J Louviere and George Woodworth. Design and analysis of simulated consumer choice or allocation experiments: an approach based on aggregate data. *Journal of marketing research*, 1983.

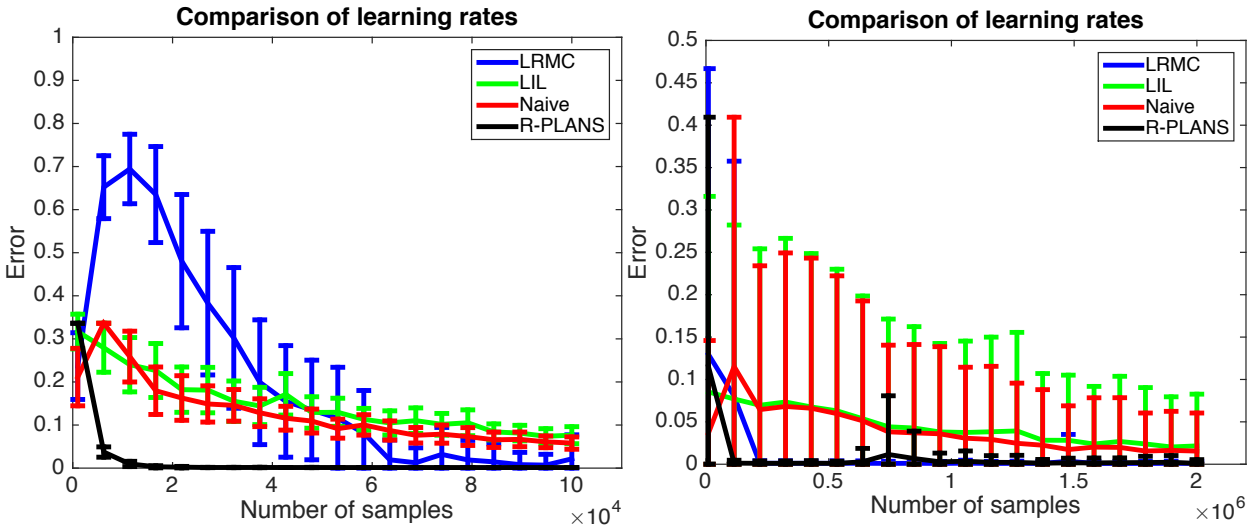

(a) ML-100K; $(200, 2)$

(b) ML-100K; $(200, 4)$

Figure 2: Error of various algorithms with increasing budget. Numbers in the brackets represent values for (K, r). The error is defined as $\boldsymbol{L}_{\hat{i},\hat{j}} - \boldsymbol{L}_{i_\star,j_\star}$ where $(\hat{i}, \hat{j})$ is a pair of optimal choices as estimated by each algorithm.

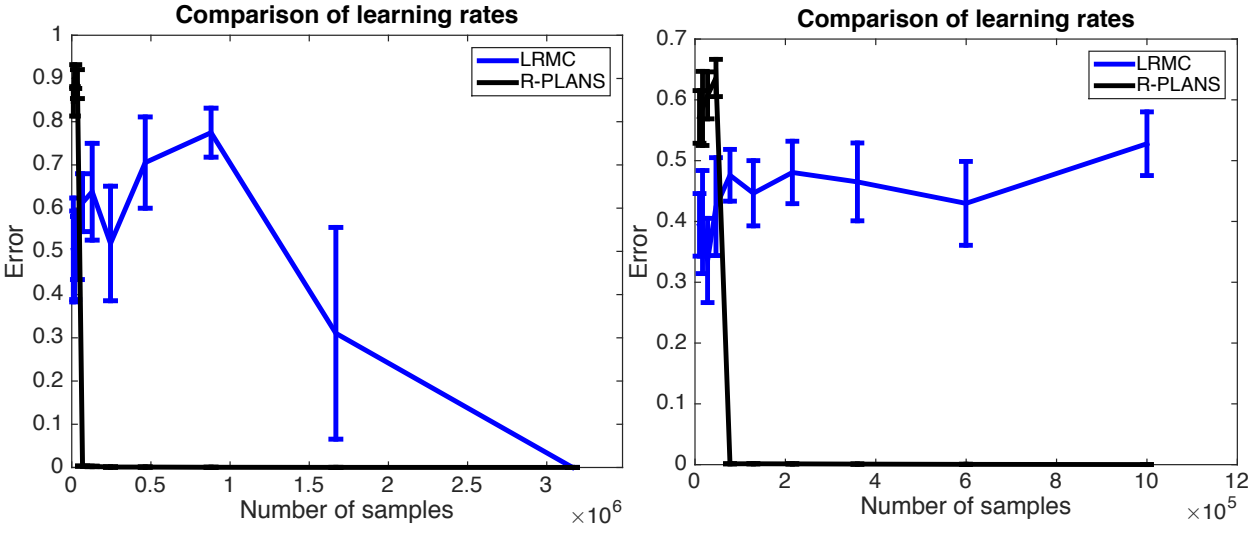

(a) ML-100K; $(200, 2)$

(b) ML-100K; $(200, 4)$

Figure 3: Error of LRMC and R-PLANS algorithms with increasing budget. Numbers in the brackets represent values for (K, r). The error is defined as $\boldsymbol{L}_{\hat{i},\hat{j}} - \boldsymbol{L}_{i_\star,j_\star}$ where $(\hat{i}, \hat{j})$ is a pair of optimal choices as estimated by each algorithm.. This is for the ML-100K dataset

[4] Emmanuel J Candès and Benjamin Recht. Exact matrix completion via convex optimization. *FOCM*, 2009.

[5] Jean-Yves Audibert and Sébastien Bubeck. Best arm identification in multi-armed bandits. In *COLT*, 2010.

[6] Victor Gabillon, Mohammad Ghavamzadeh, Alessandro Lazaric, and Sébastien Bubeck. Multi-bandit best arm identification. In *NIPS*, 2011.

[7] Sébastien Bubeck, Tengyao Wang, and Nitin Viswanathan. Multiple identifications in multi-armed bandits. In *ICML*, 2013.

[8] Eyal Even-Dar, Shie Mannor, and Yishay Mansour. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *JMLR*, 2006.

[9] Satyen Kale, Lev Reyzin, and Robert E Schapire. Non-stochastic bandit slate problems. In *NIPS*, 2010.

[10] Filip Radlinski, Robert Kleinberg, and Thorsten Joachims. Learning diverse rankings with multi-armed bandits. In *ICML*. ACM, 2008.

[11] Matthew Streeter and Daniel Golovin. An online algorithm for maximizing submodular functions. In *NIPS*, 2009.

[12] Yisong Yue and Carlos Guestrin. Linear submodular bandits and their application to diversified retrieval. In *NIPS*, pages 2483–2491, 2011.

[13] Max Simchowitz, Kevin Jamieson, and Ben Recht. Best-of-k bandits. In *COLT*, 2016.

[14] Fan Guo, Chao Liu, Anitha Kannan, Tom Minka, Michael Taylor, Yi-Min Wang, and Christos Faloutsos. Click chain model in web search. In *WWW*. ACM, 2009.

[15] Nick Craswell, Onno Zoeter, Michael Taylor, and Bill Ramsey. An experimental comparison of click position-bias models. In *WSDM*. ACM, 2008.

[16] Branislav Kveton, Csaba Szepesvari, Zheng Wen, and Azin Ashkan. Cascading bandits: Learning to rank in the cascade model. In *ICML*, pages 767–776, 2015.

[17] Jaya Kawale, Hung H Bui, Branislav Kveton, Long Tran-Thanh, and Sanjay Chawla. Efficient thompson sampling for online matrix-factorization recommendation. In *NIPS*, 2015.

[18] William E Bishop and M Yu Byron. Deterministic symmetric positive semidefinite matrix completion. In *Advances in Neural Information Processing Systems*, pages 2762–2770, 2014.

[19] Alex Gittens and Michael Mahoney. Revisiting the nystrom method for improved large-scale machine learning. In *ICML*, pages 567–575, 2013.

[20] Petros Drineas and Michael W Mahoney. On the nyström method for approximating a gram matrix for improved kernel-based learning. *JMLR*, 2005.

[21] Amit Deshpande, Luis Rademacher, Santosh Vempala, and Grant Wang. Matrix approximation and projective clustering via volume sampling. In *SODA*, 2006.

[22] Sanjiv Kumar, Mehryar Mohri, and Ameet Talwalkar. Sampling methods for the nyström method. *JMLR*, 2012.

[23] Kevin Jamieson, Matthew Malloy, Robert Nowak, and Sébastien Bubeck. lil'ucb: An optimal exploration algorithm for multi-armed bandits. *COLT*, 2014.

[24] Raghunandan Keshavan, Andrea Montanari, and Sewoong Oh. Matrix completion from noisy entries. In *Advances in Neural Information Processing Systems*, pages 952–960, 2009.

[25] F Maxwell Harper and Joseph A Konstan. The movielens datasets: History and context. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 5(4):19, 2015.

[26] Joel A Tropp. An introduction to matrix concentration inequalities. *arXiv preprint arXiv:1501.01571*, 2015.